CRAY

Cray® Urika-GD™

Technical Specifications

The Urika-GD™ graph analytics appliance designed for data discovery in very large datasets using graph analytics. The Urika appliance is available in a range of different sizes. Contact Cray for specifications of appliances larger than 512 processors (Urika-512).

The Urika-GD appliance consists of:

• Graph analytics platform, providing graph-optimized hardware with shared-memory, multithreading and scalable I/O
• Graph analytics database, providing an RDF triplestore and SPARQL query engine
• Graph analytics application services, providing management, security and data pipeline functions

GRAPH ANALYTICS PLATFORM	Urika-64	Urika-128	Urika-256	Urika-512
Cabinets - Processor	1	2	3	6
Cabinets — Storage	1	1	1	1
Processors				
Threadstorm4 Graph Accelerators with128 hardware threads per processor	64	128	256	512
x86 Management and I/O	34	68	34	68
Global Shared Memory (TB)	2	4	8	16
External Connectivity				
10 Gigabit Ethernet Ports	1	2	2	2
Gigabit Ethernet ports	8	8	8	8
Infiniband HCAs	Optional	Optional	Optional	Optional
Fibre Channel HBAs	4	4	4	4

Interconnect	Three-dimensional torus interconnect using the Cray Seastar2 communications processor. Each SeaStar2 provides six 7.6 GB/s links to neighbors, and a 6.4 GB/s connection to the x86 or Threadstorm processor. Total interconnect bandwidth scales linearly with increasing system size.
Storage	Includes 40 TB boot RAID with preinstalled system software. Natively supports POSIX-compliant, Lustre® parallel file system for scalable I/O performance on user storage. Optional support for GPFS, Panassas parallel file systems. NFS supported for access to other file systems
Cabinets – Storage	Includes 40 TB boot RAID with pre-installed system software. Natively supports POSIX-compliant, Lustre Parallel Filesystem for scalable I/O performance on user storage. Optional support for GPFS, Panassas parallel file systems. NFS supported for access to other file systems.
Cabinet Dimensions	Processor cabinets: 80.5" high x 22.5" wide x 56.75" deep (2,045mm H x 572mm W x 1,441mm D) Storage cabinets: 78" high x 23.75" wide x 41" deep (1,981mm H x 603mm W x 1,041mm D)
Weight	1,600 lbs. (726 kg) per processor cabinet, 1,500 lbs. (680 kg) per storage cabinet.

GRAPH ANALYTICS PLATFORM CONTINUED

Voltage	208 VAC2, 3 phase, 60 or 50 Hz for processor cabinets 200/208 V AC single phase for storage cabinets
Power	20.0 kW (20.41kVA)/ processor cabinet maximum. 11.4 kW (11.95 kVA)/storage cabinet, maximum.
Cooling	Air cooled Processor cabinets: maximum airflow 3,000 cfm, intake on the bottom, exhaust to top. Storage cabinets: maximum airflow 800 cfm, airflow from front to back.
Reliability	Hot-swappable power supplies, real-time fault monitoring, software failover management forcritical functions, dedicated System Management Workstation
Operating Conditions	Ambient air: 60-80°F (16-27°C); under-floor air: 50-60°F (10-16°C); relative humidity: 30-65% noncondensing; altitude sea level to 10,000 feet (0 to 3,048m)
Certificates	Safety: CE (2006/95/EC), UL 60950-1 (2nd Ed.); CSA C22.2 No. 60950-1-07 EMC: CE (2004/108/EC), ICES-0003, FCC Part 15, Subpart B; BSMI, VCCI, C-Tick Other: RoHS 2 (2011/65/EU)

GRAPH ANALYTICS DATABASE

Data Store	In-memory W3C compliant RDF triplestore, with an optional graph name associated with each triple, enabling simultaneous loading of multiple disparate graphs [thus, the data store is really a "quad store"].
Data Query	W3C standard SPARQL 1.1 query. Query transport via SPARQL over HTTP
Data Ingest	Batch load of RDF triple store from Lustre file system Dynamic, in-memory database updates.
Inferencing	Forward chaining inferencing from rule-sets
Result Sets	Query results returned via network API Large query result sets optionally returned via high-performance Lustre file system
External Integration	Integrates with a range of industry-leading data integration, visualization and business intelligence environments. Contact Cray for the current list.

GRAPH ANALYTICS APPLICATION SERVICES

Interface	Web-based GUI and command line interface can be used for data and database management Queries and query results transported over HTTP
Appliance Management	Data and database management, performance and monitoring dashboard, inferencing rule management, query queue management, role-based user management
Checkpoints	Periodic and application driven checkpoint of in-memory graph
Logging	Appliance and database logs

1 External connectivity (the number of 10 Gigabit Ethernet, Gigabit Ethernet, Infiniband and Fibre Channel ports) can be optionally modified to suit customer needs. Standard configurations are shown.
2 480 VAC three-phase power option is also available if preferred.

About Urika-GD The Urika-GD big data appliance for graph analytics helps enterprises gain key insights by discovering relationships in big data. Its highly scalable, real-time graph analytics warehouse supports ad hoc queries, pattern-based searches, inferencing and deduction. The Urika-GD appliance complements an existing data warehouse or Hadoop® cluster by offloading graph workloads and interoperating within the existing analytics workflow.

About Cray Global supercomputing leader Cray Inc. provides innovative systems and solutions enabling scientists and engineers in industry, academia and government to meet existing and future simulation and analytics challenges. Leveraging more than 40 years of experience in developing and servicing the world's most advanced supercomputers, Cray offers a comprehensive portfolio of supercomputers and big data storage and analytics solutions delivering unrivaled performance, efficiency and scalability. Go to www.cray.com for more information.